Two Notions of Fairness in Random Projections

Random projections are often discussed in terms of distance preservation or computational efficiency. An equally interesting perspective is **what kind of fairness the projection is designed to preserve**. Two common constructions illustrate an important trade-off.

HD projections: fairness of information distribution

For a structured projection such as

$$y = HDx,$$

where H is a Hadamard matrix and D is a random diagonal sign matrix:

- Every input coordinate contributes to every output.
- Every contribution has exactly the same magnitude.
- Only the signs differ.
- Information is spread uniformly across the representation.

In this sense, every feature receives **equal representation**.

If the outputs are binarized, each random hyperplane has a normal vector of the form

$$(\pm 1, \pm 1, \ldots, \pm 1),$$

so every coordinate participates equally in every decision boundary.

Gaussian random projections: fairness of response directions

For a dense Gaussian random projection:

- Hyperplane normals are uniformly distributed over the sphere.
- No geometric direction is preferred.
- Individual coefficients vary in magnitude.
- Some outputs are naturally more sensitive to particular input coordinates than others.

Here, the projection is **rotationally invariant**, rather than coordinate invariant.

Two different symmetries

These constructions optimize different notions of fairness.

HD projections

- Equal influence for every coordinate.
- Uniform information spreading.

- Hypercube symmetry (coordinate permutations and sign flips).

Gaussian projections

- Equal treatment of every direction in space.
- Isotropic geometric response.
- Full rotational symmetry.

These are distinct objectives, and neither implies the other.

A useful way to think about it

The trade-off can be summarized as:

> **Maximum fairness of information distribution** versus **maximum fairness of response directions.**

HD projections ask every feature to "vote" with equal strength.

Gaussian projections allow unequal votes in order to sample directions uniformly.

Neither approach is universally superior—they simply encode different inductive biases.

Why this matters

For applications such as locality-sensitive hashing, structured neural networks, or random feature generation, the choice affects more than computational efficiency.

It determines whether the projection is designed primarily to:

- distribute information as evenly as possible among features, or
- probe the geometry of the space as uniformly as possible.

This is a subtle distinction that is easy to overlook.

Similar trade-offs appear elsewhere

Once recognized, this pattern appears in many disciplines.

- **Crystals vs. liquids:** crystals distribute atoms in highly regular arrangements but exhibit directional anisotropy, while liquids are statistically isotropic.
- **Digital images:** square pixel grids treat every pixel identically, yet horizontal and diagonal directions are not equivalent.
- **Numerical simulation:** Cartesian grids provide uniform computation but introduce directional discretization errors.

- **Signal processing:** structured transforms such as the Hadamard transform emphasize regularity and balanced energy propagation, whereas random transforms emphasize isotropic sampling.
- **Neural network initialization:** orthogonal and Hadamard-based initializations favor balanced propagation of information, while Gaussian initializations emphasize isotropic statistics.

Final thought

One way to view the design space is that **every random projection chooses which symmetry it wishes to preserve**.

Structured Hadamard-based projections preserve fairness among coordinates.

Gaussian projections preserve fairness among geometric directions.

Recognizing this distinction helps clarify why different random projection schemes can behave differently even when they have similar theoretical guarantees for distance preservation. It also suggests that the "best" projection depends on the role the projection plays in the overall system, rather than there being a universally optimal choice.